

Amdt. dated August 18, 2006
Reply to Office action of June 12, 2006

Serial No. 10/721,018
Docket No. TUC920030149US1
Firm No. 0022.0067

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims

- 1 1. (original) A method comprising:
 - 2 maintaining in a first data structure in a first storage controller, a first index for a
 - 3 first write data task for writing data to a storage device coupled to the first storage
 - 4 controller and for writing data to a storage device coupled to a second storage controller,
 - 5 wherein the first write data task comprises a first sequence of data writing subtasks in
 - 6 which each subtask has a sequence number identifying the position of the subtask in the
 - 7 sequence of subtasks of the first write data task and wherein the first index identifies the
 - 8 sequence number of the next subtask in sequence to be sent by a first processor of the
 - 9 first storage controller to a second processor of the first storage controller;
 - 10 adding a subtask of said first sequence of data writing subtasks to a queue;
 - 11 comparing the sequence number of a subtask in the queue to the index of the first
 - 12 data structure; and
 - 13 if the subtask in the queue has the sequence number identified by the index,
 - 14 sending the subtask to a second processor of the first storage controller to generate a
 - 15 write command to the second storage controller.
- 1 2. (original) The method of claim 1 further comprising:
 - 2 maintaining in a second data structure in the first storage controller, a second
 - 3 index for a second write data task for writing data to a storage device coupled to the first
 - 4 storage controller and for writing data to a storage device coupled to the second storage
 - 5 controller, wherein the second write data task comprises a second sequence of data
 - 6 writing subtasks in which each subtask has a sequence number identifying the position of
 - 7 the subtask in the sequence of subtasks of the second write data task and wherein the

8 second index identifies the sequence number of the next subtask in sequence of the
9 second write data task to be sent to a second processor of the first storage controller.

10

1 3. (original) The method of claim 2 further comprising:
2 adding a subtask of said second sequence of data writing subtasks to said queue;
3 comparing the sequence number of a subtask in the queue to the index of the
4 second data structure; and
5 if the subtask in the queue has the sequence number identified by the index of the
6 second data structure, sending the subtask to a second processor of the first storage
7 controller to generate a write command to the second storage controller.

1 4. (original) The method of claim 3 wherein each subtask has a field identifying
2 the write data task of which it is a subtask and wherein each data structure has a field
3 identifying the write data task for which the index of the data structure identifies the
4 sequence number of the next subtask in sequence for the write data task of the data
5 structure, said method further comprising:

6 identifying the data structure containing the index to be used for comparing the
7 sequence number of a subtask in the queue to the index of the identified data structure,
8 using the subtask field identifying the write data task of which it is a subtask and using
9 the data structure field identifying the write data task of the data structure.

1 5. (original) The method of claim 4 wherein said adding a subtask of said first
2 sequence of data writing subtasks to said queue includes adding a subtask of said first
3 sequence of data writing subtasks to a first subqueue of said queue and wherein said
4 adding a subtask of said second sequence of data writing subtasks to said queue includes
5 adding a subtask of said second sequence of data writing subtasks to a second subqueue
6 of said queue.

Amdt. dated August 18, 2006
Reply to Office action of June 12, 2006

Serial No. 10/721,018
Docket No. TUC920030149US1
Firm No. 0022.0067

1 6. (original) The method of claim 5 wherein said first subqueue of subtasks of
2 said first sequence of data writing subtasks is sorted in sequential order and wherein said
3 second subqueue of subtasks of said second sequence of data writing subtasks is sorted in
4 sequential order.

1 7. (original) The method of claim 6 wherein said identifying and comparing is
2 repeated for each subtask in said queue each time a subtask is added to the queue.

1 8. (original) The method of claim 7 wherein the sequence number of each
2 subtask corresponds to a track of a storage device coupled to said second storage
3 controller.

1 9. (currently amended) An article of manufacture wherein the article of
2 manufacture A storage medium having at least one of code and logic implemented
3 therein which causes operations, the operations comprising:

4 maintaining in a first data structure in a first storage controller, a first index for a
5 first write data task for writing data to a storage device coupled to the first storage
6 controller and for writing data to a storage device coupled to a second storage controller,
7 wherein the first write data task comprises a first sequence of data writing subtasks in
8 which each subtask has a sequence number identifying the position of the subtask in the
9 sequence of subtasks of the first write data task and wherein the first index identifies the
10 sequence number of the next subtask in sequence to be sent by a processor of the first
11 storage controller to a second processor of the first storage controller;

12 adding a subtask of said first sequence of data writing subtasks to a queue;
13 comparing the sequence number of a subtask in the queue to the index of the first
14 data structure; and

Amdt. dated August 18, 2006
Reply to Office action of June 12, 2006

Serial No. 10/721,018
Docket No. TUC920030149US1
Firm No. 0022.0067

15 if the subtask in the queue has the sequence number identified by the index,
16 sending the subtask to a second processor of the first storage controller to generate a
17 write command to the second storage controller.

1 10. (currently amended) The storage medium article of claim 9, the
2 operations further comprising:
3 maintaining in a second data structure in the first storage controller, a second
4 index for a second write data task for writing data to a storage device coupled to the first
5 storage controller and for writing data to a storage device coupled to the second storage
6 controller, wherein the second write data task comprises a second sequence of data
7 writing subtasks in which each subtask has a sequence number identifying the position of
8 the subtask in the sequence of subtasks of the second write data task and wherein the
9 second index identifies the sequence number of the next subtask in sequence of the
10 second write data task to be sent by a processor of the first storage controller to a second
11 processor of the first storage controller.

1 11. (currently amended) The storage medium article of claim 10, the
2 operations further comprising:
3 adding a subtask of said second sequence of data writing subtasks to said queue;
4 comparing the sequence number of a subtask in the queue to the index of the
5 second data structure; and
6 if the subtask in the queue has the sequence number identified by the index of the
7 second data structure, sending the subtask to a second processor of the first storage
8 controller to generate a write command to the second storage controller.

1 12. (currently amended) The storage medium article of claim 11 wherein
2 each subtask has a field identifying the write data task of which it is a subtask and
3 wherein each data structure has a field identifying the write data task for which the index

Amdt. dated August 18, 2006
Reply to Office action of June 12, 2006

Serial No. 10/721,018
Docket No. TUC920030149US1
Firm No. 0022.0067

4 of the data structure identifies the sequence number of the next subtask in sequence for
5 the write data task of the data structure, the operations further comprising:

6 identifying the data structure containing the index to be used for comparing the
7 sequence number of a subtask in the queue to the index of the identified data structure,
8 using the subtask field identifying the write data task of which it is a subtask and using
9 the data structure field identifying the write data task of the data structure.

1 13. (currently amended) The storage medium article of claim 12 wherein
2 said adding a subtask of said first sequence of data writing subtasks to said queue
3 includes adding a subtask of said first sequence of data writing subtasks to a first
4 subqueue of said queue and wherein said adding a subtask of said second sequence of
5 data writing subtasks to said queue includes adding a subtask of said second sequence of
6 data writing subtasks to a second subqueue of said queue.

1 14. (currently amended) The storage medium article of claim 13 wherein
2 said first subqueue of subtasks of said first sequence of data writing subtasks is sorted in
3 sequential order and wherein said second subqueue of subtasks of said second sequence
4 of data writing subtasks is sorted in sequential order.

1 15. (currently amended) The storage medium article of claim 14 wherein
2 said identifying and comparing is repeated for each subtask in said queue each time a
3 subtask is added to the queue.

1 16. (currently amended) The storage medium article of claim 15 wherein the
2 sequence number of each subtask corresponds to a track of a storage device coupled to
3 said second storage controller.

Amdt. dated August 18, 2006
Reply to Office action of June 12, 2006

Serial No. 10/721,018
Docket No. TUC920030149US1
Firm No. 0022.0067

- 1 17. (original) A system for use with a remote storage controller having a
2 storage device coupled to said remote storage controller, comprising:
3 a first storage controller including a plurality of processors, a first data structure, a
4 second data structure, and a queue; said first storage controller having a storage device
5 coupled to said first storage controller;
6 means for maintaining in said first data structure, a first index for a first write data
7 task for writing data to said storage device coupled to the first storage controller and for
8 writing data to a storage device coupled to a remote storage controller, wherein the first
9 write data task comprises a first sequence of data writing subtasks in which each subtask
10 has a sequence number identifying the position of the subtask in the sequence of subtasks
11 of the first write data task and wherein the first index identifies the sequence number of
12 the next subtask in sequence to be sent to a second processor of the first storage
13 controller;
14 means for adding a subtask of said first sequence of data writing subtasks to said
15 queue;
16 means for comparing the sequence number of a subtask in the queue to the index
17 of the first data structure; and
18 means for, if the subtask in the queue has the sequence number identified by the
19 index, sending the subtask to a second processor of the first storage controller to generate
20 a write command to the remote storage controller.

- 1 18. (original) The system of claim 17 further comprising:
2 means for maintaining in the second data structure of the first storage controller, a
3 second index for a second write data task for writing data to a storage device coupled to
4 the first storage controller and for writing data to a storage device coupled to the remote
5 storage controller, wherein the second write data task comprises a second sequence of
6 data writing subtasks in which each subtask has a sequence number identifying the
7 position of the subtask in the sequence of subtasks of the second write data task and

Amdt. dated August 18, 2006
Reply to Office action of June 12, 2006

Serial No. 10/721,018
Docket No. TUC920030149US1
Firm No. 0022.0067

8 wherein the second index identifies the sequence number of the next subtask in sequence
9 of the second write data task to be sent to a second processor of the first storage
10 controller.

1 19. (original) The system of claim 18 further comprising:
2 means for adding a subtask of said second sequence of data writing subtasks to
3 said queue;
4 means for comparing the sequence number of a subtask in the queue to the index
5 of the second data structure; and
6 means for, if the subtask in the queue has the sequence number identified by the
7 index of the second data structure, sending the subtask to a second processor of the first
8 storage controller to generate a write command to the remote storage controller.

1 20. (original) The system of claim 19 wherein each subtask has a field
2 identifying the write data task of which it is a subtask and wherein each data structure has
3 a field identifying the write data task for which the index of the data structure identifies
4 the sequence number of the next subtask in sequence for the write data task of the data
5 structure, said system further comprising:
6 means for identifying the data structure containing the index to be used for
7 comparing the sequence number of a subtask in the queue to the index of the identified
8 data structure, using the subtask field identifying the write data task of which it is a
9 subtask and using the data structure field identifying the write data task of the data
10 structure.

1 21. (original) The system of claim 20 wherein said queue includes a first
2 subqueue and a second subqueue and said means for adding a subtask of said first
3 sequence of data writing subtasks to said queue includes means for adding a subtask of
4 said first sequence of data writing subtasks to said first subqueue of said queue and

Amdt. dated August 18, 2006
Reply to Office action of June 12, 2006

Serial No. 10/721,018
Docket No. TUC920030149US1
Firm No. 0022.0067

5 wherein said means for adding a subtask of said second sequence of data writing subtasks
6 to said queue includes means for adding a subtask of said second sequence of data
7 writing subtasks to said second subqueue of said queue.

1 22. (original) The system of claim 21 wherein said means for adding a
2 subtask of said first sequence of data writing subtasks includes means for sorting said
3 first subqueue of subtasks of said first sequence of data writing subtasks in sequential
4 order and wherein means for adding a subtask of said second sequence includes means
5 for sorting said second subqueue of subtasks of said second sequence of data writing
6 subtasks in sequential order.

1 23. (original) The system of claim 22 wherein said identifying and
2 comparing is repeated for each subtask in said queue each time a subtask is added to the
3 queue.

1 24. (original) The system of claim 23 wherein the sequence number of
2 each subtask corresponds to a track of a storage device coupled to said remote storage
3 controller.